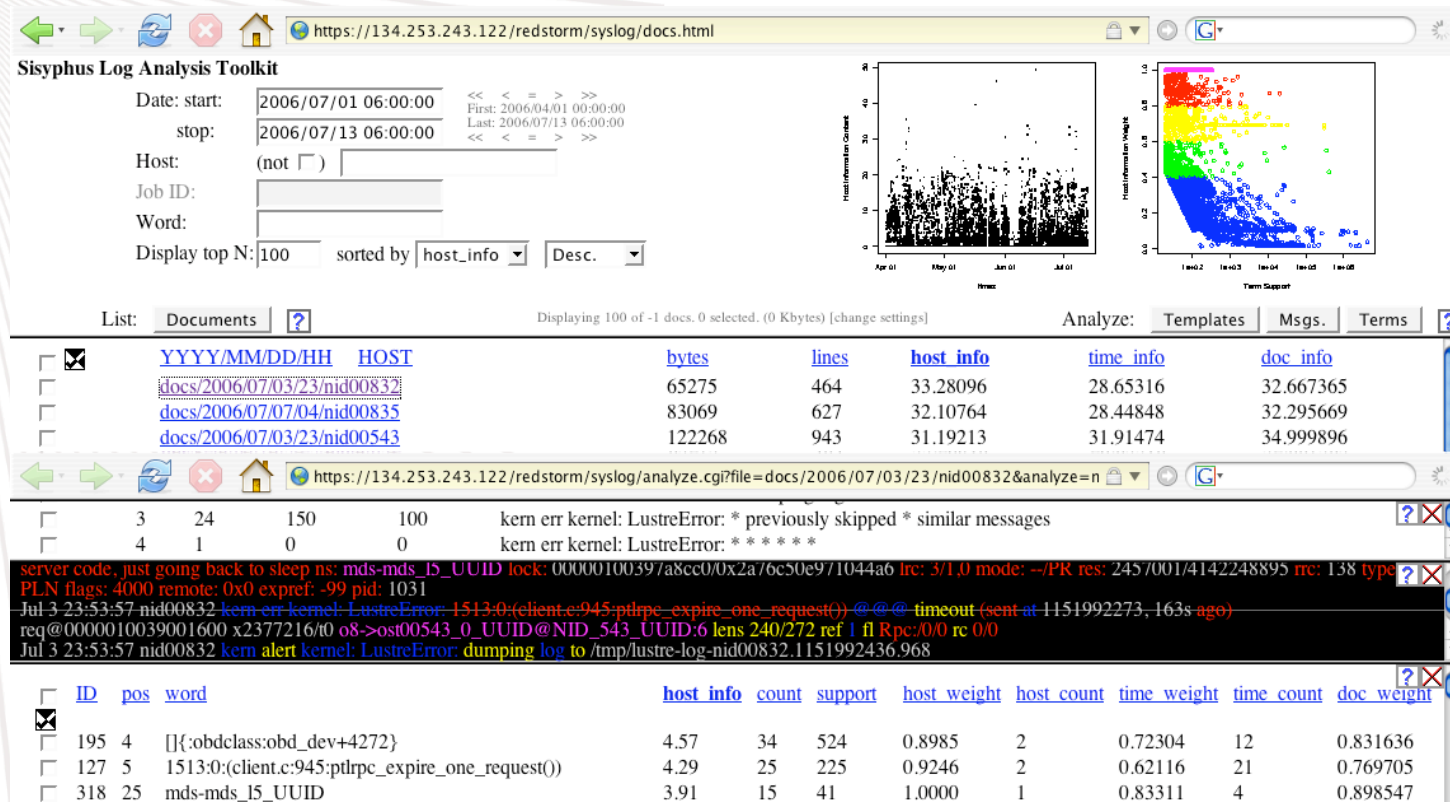




Sisyphus - Increasing RAS (Reliability, Availability, Serviceability) via Informatic Analysis of Event Logs



<http://www.cs.sandia.gov/sisyphus>
"the cure for logorrhea"

Sandia National Laboratories


Jon Stearley <jrstear@sandia.gov>



Sisyphus - Increasing RAS (Reliability, Availability, Serviceability) via Informatic Analysis of Event Logs

Event logs are:

- **Ubiquitous** (nearly all types of computing systems generate them)
- **Informational** (nearly always the first place to look for problems)
- **Voluminous** (infeasible to read all of them)
- **Noisy** (many messages are completely useless)
- **Non-stationary** (characteristics change with time)
- **Ambiguous** (meaning often depends on operational context)





Sisyphus - Increasing RAS (Reliability, Availability, Serviceability) via Informatic Analysis of Event Logs

Interrogating (grep) or monitoring (logsurfer) with regular expressions:

- **Is common**
- **Can be easy (or NOT)**
- **Can be effective (or NOT)**

Many hard questions cannot be easily answered with grep!

Under what conditions does string “xyz” occur?

^ e.g. nodes, times, apps, users...

What do you look for... when you don't know what to look for?

Is anything “strange” in the logs?



Sisyphus - Increasing RAS (Reliability, Availability, Serviceability) via Informatic Analysis of Event Logs

(1) Sisyphus Log Analysis Toolkit

Date: start: 2006/07/01 06:00:00 First: 2006/04/01 00:00:00
stop: 2006/07/13 06:00:00 Last: 2006/07/13 06:00:00
Host: (not ☐)
Job ID:
Word:
Display top N: 100 sorted by host_info Desc.

(2) List: Documents Analyze: Templates Msgs. Terms

	YYYY/MM/DD/HH	HOST	bytes	lines	host_info	time_info	doc_info
(2)	docs/2006/07/03/23/nid00832		65275	464	33.28096	28.65316	32.667365
	docs/2006/07/07/04/nid00835		83069	627	32.10764	28.44848	32.295669
	docs/2006/07/03/23/nid00543		122268	943	31.19213	31.91474	34.999896

(4) 3 24 150 100 kern err kernel: LustreError: * previously skipped * similar messages
4 1 0 0 kern err kernel: LustreError: *****

(5) server code, just going back to sleep ns: mds-mds_l5_UUID lock: 00000100397a8cc0/0x2a/6c50e9/1044a6 lrc: 5/1/0 mode: --PR res: 245/001/4142248895 rc: 138 type
PLN flags: 4000 remote: 0x0 expref: -99 pid: 1031
Jul 3 23:53:57 nid00832 kern err kernel: LustreError: 1513:0:(client.c:945:ptlrpc_expire_one_request()) @@@@ timeout (sent at 1151992273.163s ago)
req@0000010039001600 x2377216/t0 o8->ost00543_0_UUID@NID_543 UUID:6 lens 240/272 ref 1 fl Rpc:0/0 rc 0/0
Jul 3 23:53:57 nid00832 kern alert kernel: LustreError: dumping log to /tmp/lustre-log-nid00832.1151992436.968

(6)

ID	pos	word	host_info	count	support	host_weight	host_count	time_weight	time_count	doc_weight	
(6)	195	4	[{obdclass:obd_dev+4272}	4.57	34	524	0.8985	2	0.72304	12	0.831636
	127	5	1513:0:(client.c:945:ptlrpc_expire_one_request())	4.29	25	225	0.9246	2	0.62116	21	0.769705
	318	25	mds-mds_l5_UUID	3.91	15	41	1.0000	1	0.83311	4	0.898547

(1) Query

(2) Document statistics

(3) Term information weights (and colors)

(4) Message templates (co-occurring terms)

(5) Information-colored logs

(6) Term statistics

Sandia National Laboratories

<http://www.cs.sandia.gov/sisyphus>

Jon Stearley <jrstear@sandia.gov>